arXiv:1704.00085v1 [cs.CV] 31 Mar 2017

Optimal Reconstruction with a Small Number of Views

Cheng Peng, Volkan Isler
University of Minnesota
200 Union St. SE Keller 4-192, Minneapolis, MN, 55455, USA
{peng0175, isler}@umn.edu

Abstract

Estimating positions of world points from features observed in images is a key problem in 3D reconstruction, image mosaicking, simultaneous localization and mapping and structure from motion. We consider a special instance in which there is a dominant ground plane $\mathcal{G}$ viewed from a parallel viewing plane $\mathcal{S}$ above it. Such instances commonly arise, for example, in aerial photography.

Consider a world point $g \in \mathcal{G}$ and its worst case reconstruction uncertainty $\varepsilon(g, \mathcal{S})$ obtained by merging all *possible views of g chosen from $\mathcal{S}$. We first show that one can pick two views s_p and s_q such that the uncertainty $\varepsilon(g, \{s_p, s_q\})$ obtained using only these two views is almost as good as (i.e. within a small constant factor of) $\varepsilon(g, \mathcal{S})$. Next, we extend the result to the entire ground plane $\mathcal{G}$ and show that one can pick a small subset of $\mathcal{S}' \subseteq \mathcal{S}$ (which grows only linearly with the area of $\mathcal{G}$) and still obtain a constant factor approximation, for every point $g \in \mathcal{G}$, to the minimum worst case estimate obtained by merging all views in $\mathcal{S}$.*

Our results provide a view selection mechanism with provable performance guarantees which can drastically increase the speed of scene reconstruction algorithms. In addition to theoretical results, we demonstrate their effectiveness in an application where aerial imagery is used for monitoring farms and orchards.

1. Introduction

Consider a scenario where a plane flying at a fixed altitude is capturing images of a ground plane below so as to reconstruct the scene (Figure 1). Over the course of its flight, the plane may capture thousands of images which can easily overwhelm image reconstruction algorithms. Our goal in this paper is to answer the question of whether we can select a small number of images and focus only on them without giving up on the reconstruction quality.

We first study a basic version where we focus on a single world point. The goal is to select a small number of images from which the 3D position of the world point can be accurately estimated (Problem 1). We then present a general version where the goal is to minimize the error for the entire scene (Problem 2) from a small set of images. Note that in the latter case, the same set of images must be used for every scene point.

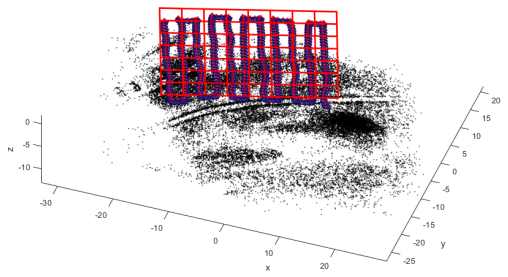

Figure 1: Reconstruction using 414 images obtained from a constant altitude. We will show that it is possible to obtain the same reconstruction quality with a much smaller number of images.

In order to formalize these two problems, we first need to formalize the error model and the uncertainty objective. Let g be a world point and I be an image taken from a camera at position s and orientation θ. Let p the observed projection of g onto I and p^* be the unobserved true projection represented as vectors originating from the camera center s. We will employ a *bounded uncertainty model* where we will assume that the angle between p and p^* is bounded by a known (or desired) quantity α. Therefore, the 3D location of the world point g is contained inside a cone C apexed at s and with symmetry axis along p and cone angle 2α.

Merging measurements: In order to estimate the true location of a world point from multiple measurements, we simply intersect the corresponding cones. The diameter of the intersection is used as an uncertainty measure. We chose diameter over the volume so as to avoid degenerate cases where the intersection has almost zero volume but large diameter which could still generate large reprojection error.

Uncertainty as worst-case reconstruction error: Rather

than associating as single cone for a specific measurement, our formulation considers a possibly infinite set of viable cones for a given true camera pose and world point pair. To do this, we consider all possible perturbations of relevant quantities (projection, location or pose). When merging measurements, we consider the worst-case scenario which maximizes the reconstruction uncertainty. This formulation gives us a deterministic worst-case error model. It also allows us to factor out unknown or uncontrollable quantities such as camera orientation. The process can be visualized as a game, where for each true projection p^* an adversary chooses a measurement p in such a way that when the cones corresponding to those measurements are merged, the uncertainty is maximized. The same reasoning applies for orientation uncertainty as well. We can treat the orientation as an uncontrolled external disturbance. The resulting estimation process can again be visualized as a new game where we choose the camera center locations but an adversary chooses the rotation of each camera and the observed projection of the feature so as to maximize the intersection area. Under a spherical camera model, these two choices can simply be added. Therefore the adversary can choose a single parameter. Of course, the choice must still ensure that the cone contains the true world point.

Let us make these concepts more concrete by revisiting the example of a plane flying at a fixed altitude. Consider the ground plane $\mathcal{G}$, a point $g \in \mathcal{G}$ and the viewing plane $\mathcal{V}$ assumed to be parallel to $\mathcal{G}$. Suppose we take every possible measurement from $\mathcal{V}$ by choosing every point as a camera center. The adversary chooses the camera the measurements and orientations, we get corresponding cones, intersect them and obtain the diameter $\varepsilon_\infty(g)$: the minimum worst case uncertainty achievable to reconstruct g under this model. The subscript is chosen to emphasize that this quantity is obtained by merging possibly infinitely many views. Can we select only a few images upfront and still guarantee small deviation from $\varepsilon_\infty(g)$? Can we select views so as to provide the same guarantee for every such g? In this paper, we show that for our motivating application, the answer is yes.

2. Contributions and Related Work

The importance of view selection for scene reconstruction is well established. One of the first view selection schemes for multi-view stereo is presented in [2]. The work of Maver and Bajcsy [9] and Kutulakos and Dyer [7] use contour information to choose viewing locations. A 2003 paper by Scott et al. [12] surveys view selection methods. Recently, Furukawa et al. [3] proposed a view selection scheme to enable large scale 3D reconstruction. Their method relies on clustering images based on overlap. The resulting optimization problem is solved iteratively. The method of Hornung et al. [5] incrementally selects images and uses a proxy to ensure coverage. Mauro et al. resort to linear programming to solve the view selection problem [8]. View selection has also been considered for image based modeling [13], object retrieval [4] and target localization [6].

In the present work, we consider a much simpler setting than the previous work: cameras on a planar view plane observing a planar world scene. We present a novel uncertainty model which allows us to characterize worst-case reconstruction error in a way that is independent of particular measurements. What differentiates our work from the previous body of work is that we present a simple view selection mechanism with theoretical performance guarantees. Specifically, we show that one can select two good views and obtain a reconstruction which is almost as good as merging all possible views from the entire viewing plane. Next, we show that a coarse grid (of resolution proportional to the scene depth) can provide a good reconstructions of the entire world plane.

Our work is also related to error analysis in stereo [11, 1]. We contribute to this line of work by analyzing the reconstruction error for two (best) cameras with respect to the reconstruction error achievable by using all possible cameras for the particular geometry we consider.

In addition to mathematical proofs, we validate our results in simulations and show how our model can handle calibration errors and uncertainty in camera pose. We also establish their practical relevance in an image mosaicking application and show that the proposed view selection methods can drastically (from an overnight computation to one that takes a few minutes) improve the speed of existing algorithms without sacrificing accuracy.

3. Problem Definition

In this section, we introduce the general sensor selection problem. Consider the world point $g \in \mathcal{G}$ and a camera (s, θ) where $s \in \mathbb{R}^3$ is the projection center and $\theta \in SO(3)$ is the orientation. Suppose we have a set of measurements $\{p_1, \ldots, p_k\}$ where each p_i is expressed as a unit vector pointing towards the observed pixel and anchored at the corresponding camera center. We need a function $f(p_1, p_2, ..., p_k) = \hat{g}$ that maps measurements to $\hat{g}$, the estimate of g. This way, we can define the estimation error to be $||g - \hat{g}||$ by choosing an error measure $|| \cdot ||$.

In this paper we will consider the following "bounded uncertainty" characterization of the error: Consider the true measurement $p^* = Proj((s, \theta), g)$ given by the projection of g onto camera (s, θ) which is also represented as a vector from s pointing toward g. We make the assumption that the angle between the measurement p and the true projection p^* is bounded by a fixed threshold α. For a given measurement p, the rays corresponding to all possible p^* corresponds to a cone denoted as $Cone_\alpha((s, \theta), p)$ as shown in Fig 2, which

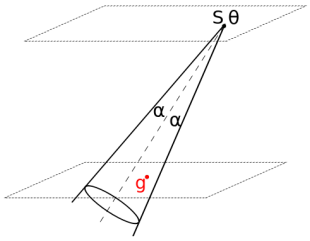

Figure 2: Right circular cone for camera (s, θ) viewing target g

is a function of both the camera parameters s and g as well as the measurement p. For the rest of the paper, we will assume a fixed α and drop the subscript. By intersecting the cones from multiple measurements p_i from views (s_i, θ_i), we can get an estimate of the true target location. The uncertainty is given by the diameter of the intersection given by $||\cap Cone((s_i, \theta_i), p_i)||$.

For sensor selection purposes, rather than a single cone, it is beneficial to associate a set of cones for each measurement. This will allow us to replace the randomness in the measurement process with a deterministic *worst-case analysis.* To do this, for a given true target location g and a camera pose (s, θ), we generate $p^* = Proj((s, \theta), g)$. Then for every possible measurement p within angle α of p^*, we define $Cone((s, \theta), p)$ and include it with the set $S(g, s, \theta)$ associated with this world point/camera pair. Note that each cone in the set includes the true location g. We can further eliminate the dependency on camera orientation by taking the union of these sets for each allowable orientation. That is, we define $S(g, s) = \bigcup_\theta S(g, s, \theta)$ with the additional requirement that $g \in Cone((s, \theta), p)$ for each cone included in the union.

We can now define the worst case uncertainty for a given set $\mathcal{S} = \{s_1, s_2, ..., s_k\}$ of camera centers and a ground point g as:

$$\varepsilon(g, \mathcal{S}) = \max_{Cone_1 \in S(g,s_1),...,Cone_k \in S(g,s_k)} ||\cap Cone_i||$$

In other words, for each camera location s_i, a cone is chosen such that the chosen cones *jointly* maximize the intersection diameter. The advantage of this formulation is that since the computation of $\varepsilon(g, \mathcal{S})$ implicitly generates all possible measurements for a given camera location and world point, it generates a worst case uncertainty independent of specific measurements and camera rotations. We are now ready to define the first problem.

Problem 1. *For a given world point g, the set of all possible viewpoints $\mathcal{S}$, a projection error bound α, and an error tolerance parameter $\rho \in \mathbb{R}$, choose a minimum cardinality subset $\mathcal{S}' \subseteq \mathcal{S}$, such that*

$$\varepsilon(g, \mathcal{S}') \leq \rho\varepsilon(g, \mathcal{S})$$

In Problem 1 the goal is to choose a small subset of camera locations whose worst case uncertainty when reconstructing a given point g is at most with a factor ρ of the worst-case uncertainty of the entire viewing set. Problem 2 generalizes it to multiple points.

Problem 2. *For a set of points $G \subseteq \mathcal{G}$, the set of all possible viewpoints $\mathcal{S}$, a projection error bound α, and an error tolerance parameter $\phi \in \mathbb{R}$, choose a minimum cardinality subset $\mathcal{S}' \subseteq \mathcal{S}$, such that*

$$\max_{g \in G} \varepsilon(g, \mathcal{S}') \leq \phi \max_{g \in \mathcal{G}} \varepsilon(g, \mathcal{S})$$

In this paper, we study a specific geometric instance of these problems where $\mathcal{G}$ and $\mathcal{S}$ are two parallel planes with distance h apart. For a given $g \in \mathcal{G}$, we will define $\varepsilon_\infty(g) = \varepsilon(g, \mathcal{S})$.

4. Sensor Selection for a Single Point

In this section, we study Problem 1 where the goal is to choose cameras to reconstruct a single point.

We will start with the two dimensional (2D) case where the ground and viewing planes reduce to lines, and the uncertainty cones become wedges. Our key result in this section is that for any point g, one can choose two cameras whose worst case uncertainty $\varepsilon_2(g)$ is almost as good as $\varepsilon_\infty(g)$ – the worst case uncertainty obtained by merging the views from *all* cameras. The key ideas in obtaining this result are: (1) if we choose two cameras at locations p and q who view g symmetrically at 90 degrees (i.e. $\angle pgq = \pi/2$), the diagonals of the worst-case uncertainty polygon (the intersection of the two wedges) are roughly of equal length. (2) Any other camera added to the sensor set can be rotated to contain the horizontal diagonal. Therefore, it does not reduce the uncertainty drastically.

4.1. The Solution of Problem 1 in 2D

Let $A = \arg\max(\varepsilon_\infty(g))$ be the set of wedges which yield the minimum worst case uncertainty. For every point c on the viewing plane, there is a wedge in A which (i) is apexed at c, (ii) has wedge angle α and (iii) contains g. By definition of $\varepsilon_\infty(g)$, the wedges are rotated so as to maximize the diagonal of the intersection.

Theorem 4.1. *Consider a target g on line G and viewing set S composed of all camera locations on a line S parallel*

to G. There exist two cameras s_p and s_q which guarantee that

$$\varepsilon_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}}\varepsilon_\infty \tag{1}$$

where $\varepsilon_\infty = \varepsilon(g, S)$ is the minimum worst case uncertainty of the entire viewing set, and ε_2 is the worst case uncertainty of $\{s_p, s_q\}$ and $0 \leq \alpha < 1/4$ is the error threshold measured in radians.

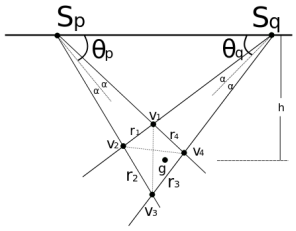

Figure 3: Notation for the two camera selection s_p and s_q

We will prove the theorem directly by providing the two cameras, computing their worst-case uncertainty ε_2 and comparing it with the minimum possible worst-case uncertainty. First, we present the notation and the setup used in the computations. We set a coordinate system whose origin is the target $g = [0, 0]$. The x-axis is on G and the z-axis points "up" toward the viewing plane. The locations of the two cameras are chosen as: $s_p = [-t/2, h]$ and $s_q = [t/2, h]$ where $t = \frac{2h}{\tan(\pi/4-\alpha)}$ and the cone orientations θ_p, θ_q respectively as shown in Fig 3. We use the angle θ between the bisector of a wedge with respect to S for orientation. Of the two half-planes whose intersection yields the wedge, the inner half plane is the one that is closer to S – i.e. the angle measured is smaller while the other half-plane is the outer half-plane also shown in Fig 3. Note that $\theta_p, \theta_q \in [\pi/4 - 2\alpha, \pi/4]$.

Their worst case uncertainty is given by

$$\varepsilon_2 = \max_{\theta_p, \theta_q} ||Cone((s_p, \theta_p), g) \cap Cone((s_q, \theta_q), g)|| \tag{2}$$

Consider the two wedges which give the worst case uncertainty (i.e. $\arg\max$ of ε_2). Let Q_{pq} be their intersection with vertices $\{v_1, v_2, v_3, v_4\}$ and edges $\{e_1, e_2, e_3, e_4\}$ as shown in Fig 3. The lengths of the edges are denoted as $r_i = Length(e_i)$ and the length of the diagonals are denoted by $diag_1, diag_2$.

We now compute these quantities.

4.1.1 Computing ε_2

In order to maximize over the orientation, we first establish the closed form solution for the edges and diagonals as functions of h,t,$\theta_{p,q}$,and α.

Using the law of sines over the triangle $s_p v_1 v_2$, we get $\frac{r_1}{\sin(2\alpha)} = \frac{\overline{s_p v_1}}{\sin \angle s_p v_2 s_q}$. We also have $\angle s_p v_2 s_q = \pi - 2\alpha - \angle s_p v_1 v_2 = \pi - 2\alpha - (\theta_p + \theta_q - 2\alpha) = \pi - \theta_p - \theta_q$. From $\triangle(s_p v_1 s_q)$, we know that $\frac{\overline{s_p v_1}}{\sin(\theta_q - \alpha)} = \frac{t}{sin(\pi - \theta_p - \theta_q + 2\alpha)}$. By combining both equations, we obtain:

$$r_1 = \frac{t \sin(\theta_q - \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q - 2\alpha)\sin(\theta_p + \theta_q)} \tag{3}$$

Using the same method, we have:

$$r_2 = \frac{t \sin(\theta_p + \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q)\sin(\theta_p + \theta_q + 2\alpha)} \tag{4}$$

From $\triangle(s_p v_3 v_4)$, we get:

$$r_3 = \frac{t \sin(\theta_q + \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q)\sin(\theta_p + \theta_q + 2\alpha)} \tag{5}$$

Similarly, from $\triangle(s_q v_1 v_4)$

$$r_4 = \frac{t \sin(\theta_p - \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q - 2\alpha)\sin(\theta_p + \theta_2)} \tag{6}$$

Using and the law of cosines, $diag_1$ can be calculated as

$$diag_1 = r_1^2 + r_2^2 - 2r_1 r_2 \cos(\theta_p + \theta_q) \tag{7}$$

Similarly, the $diag_2$ can be calculated as

$$diag_2 = r_1^2 + r_4^2 - 2r_1 r_4 \cos(\pi - \theta_p - \theta_q + 2\alpha) \tag{8}$$

We now consider the vertical diagonal whose length $diag_1$ (given in Equation 7). It is maximized when $\theta_p = \theta_q = \pi/4$. Fig 4 shows $diag_1$ as a function of the two wedge angles θ_p and θ_q and for a fixed $\alpha = 0.0113$ rad. When $\theta_p = \theta_q = \pi/4$, the vertex $v_1 = g$, which means that the inner half-planes of $Cone_p$ and $Cone_q$ intersect at g.

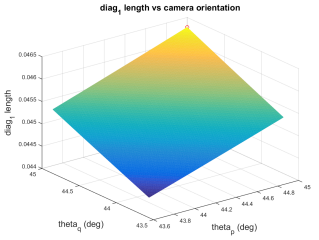

Figure 4: $diag_1$ length as a function of θ_p and θ_q

We can therefore set $\theta_p = \theta_q = \pi/4$ and write the equation of $diag_1$ as a function of α and h: Using the law of sines on the triangle $\triangle(s_q v_1 v_3)$ and $\overline{v_1 s_q} = h/\sin(\pi/4 -$

$\alpha)$, we obtain:

$$\frac{diag_1}{\sin(2\alpha)} = \frac{\overline{v_1 s_q}}{\sin(\frac{\pi}{2} - \theta - \alpha)}$$
$$diag_1 = \frac{h\sin(2\alpha)}{\cos(\theta+\alpha)\sin(\theta-\alpha)}$$
$$= \frac{h\sin(2\alpha)}{\frac{1}{2}[\sin(2\theta) - \sin(2\alpha)]}$$
$$= \frac{2h\sin(2\alpha)}{1-\sin(2\alpha)}$$

This establishes the maximum length of the diagonal $diag_1 = \frac{2h\sin(2\alpha)}{1-\sin(2\alpha)}$ in the worst case configuration of $\theta_p = \theta_q = \pi/4$.

We now compare $\varepsilon_2(s_p, s_q) = \max ||Q_{pq}||$ with ε_∞.

Lemma 4.2. *Consider the two cameras s_p, s_q in the optimal configuration described above and let $diag_1$ be the intersection of their worst-case uncertainty polygon Q_{pq}. Any cone starting from location $s_k \in A - \{s_p, s_q\}$, can be rotated to an angle θ_k such that both g and $diag_1$ are contained in its uncertainty wedge $Cone((s_k, \theta_k), g)$.*

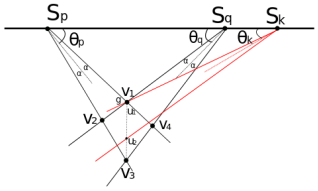

Figure 5: If the cone created by s_k that does not contain $diag_1$, we get a contradiction (proof of Lemma 4.2).

Proof. We prove the lemma by contradiction: Suppose there exists a camera $s_k \in A - \{s_p, s_q\}$ such that $Cone((s_k, \theta_k), g)$ intersects $\overline{v_1 v_3}$ at point u_1, u_2, where $u_1 \leq v_1$ and $u_2 \geq v_3$ as shown in Fig 5; Since $Cone((s_k, \theta_k), g)$ must contain target g, $u_1 = v_1$. We know that u_1, u_2 are on the vertical line passing through g, we can formulate $\overline{u_1 v_2}$ using the law of sine from the triangle $\triangle(s_k u_1 u_2)$.

$$\frac{\overline{u_1 u_2}}{\sin(2\alpha)} = \frac{h/\sin(\theta_k - \alpha)}{\sin(\pi/2 - \theta_k - \alpha)}$$
$$\overline{u_1 u_2} = \frac{h\sin(2\alpha)}{\sin(\theta_k - \alpha)\cos(\theta_k + \alpha)}$$
$$\overline{u_1 u_2} = \frac{2h\sin(2\alpha)}{\sin(2\theta_k) - \sin(2\alpha)}$$

Since $u_2 \geq v_3$, we want to find the minimum $\overline{u_1 u_2}$ by choosing different $s_k \neq s_p, s_q$, which is equivalent to minimizing $\overline{u_1 u_2}$ w.r.t. θ_k. Thus, $\overline{u_1 u_2}$ is minimized when $\sin(2\theta_k) = 1$, which results in $\theta_k = \pi/4$. By substituting $\theta_k = \pi/4$, $\overline{v_1 v_k} = \frac{2h\sin(2\alpha)}{1-\sin(2\alpha)} = diag_1$. It means that either $s_k = s_q$ or $\overline{u_1 u_2} \geq \overline{v_1 v_3}$, both of which contradict with our assumption. □

Now that we established that two cameras suffice, we compute the uncertainty value:

Lemma 4.3. *Given the two cameras s_p, s_q, the intersection polygon Q_{pq}, the maximum length of the diagonal $diag_1 = \frac{2h\sin(2\alpha)}{1-\sin(2\alpha)}$ when $\theta_p = \theta_q = \pi/4$, and the worst case uncertainty $\varepsilon_2 = \max ||Q_{pq}||$.*

$$\varepsilon_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} \cdot \frac{4h\alpha}{1-2\alpha} \tag{9}$$

The proof is by direct computation and will be given in Appendix of the supplementary material.

We conclude by presenting the proof of Theorem 4.1.

Proof. Combining Lemma 4.2 and Lemma 4.3, we can conclude that $diag_1 \leq \varepsilon_\infty \leq diag_2$. Therefore, $\varepsilon_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} \cdot \varepsilon_\infty$ □

In this section, we showed that there exist two cameras s_p and s_q with orientation $\theta_p = \theta_q = \pi/4$ such that their worst case uncertainty $\varepsilon_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} \cdot \varepsilon_\infty$. We will call the pair of cameras s_p, s_q as the **optimal pair** for the rest of the paper and this configuration as the **optimal configuration** of $\{s_p, s_q\}$.

4.2. The Solution of Problem 1 in 3D

The results of the previous section readily extend to ε_∞ in 3-D.

Theorem 4.4. *Given a target $g \in \mathcal{G}$ and a set of cameras $s \in \mathcal{S}$, where the distance between $\mathcal{G}$ and $\mathcal{S}$ is h and the number of cameras in $\mathcal{S}$ is unbounded, we claim that the optimal pair s_p and s_q gives*

$$\varepsilon_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} \cdot \varepsilon_\infty \tag{10}$$

where the minimum worst case uncertainty in 3-D is $\varepsilon_\infty = \varepsilon(g, \mathcal{S})$ and worst case uncertainty from two cameras s_p and s_q is ε_2.

To prove the theorem, all we have to do is to observe that the diagonal of a perpendicular cross section of the cone bounds the uncertainty in 3D as well. See Fig 6. Therefore, we can apply Theorem 4.1.

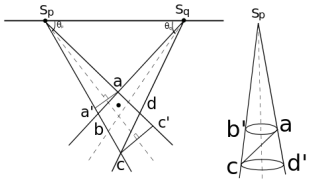

Figure 6: Uncertainty in 3D given by two intersecting cones

5. Sensor Selection for the Entire Scene

In the previous section, we established that for a world point g, the optimal pair of cameras can produce a reconstruction with approximation ratio less than $\sqrt{\frac{1+2\alpha}{1-4\alpha}}$ (Theorem 4.4). However, if we use this directly for every scene point, we may end up choosing two cameras for each scene point, which in turn might result in a large number of cameras.

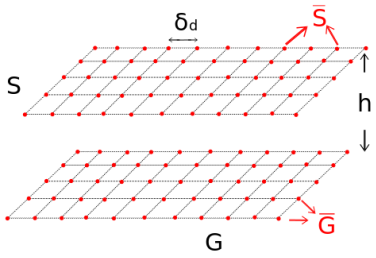

Figure 7: The square sensor grid in 3D

In this section, we show that a coarse grid of cameras provide a good reconstruction for every scene point.

Recall that $\mathcal{G}$ is the ground plane, $\mathcal{S}$ is the view plane, $\mathcal{G}$ is parallel to $\mathcal{P}$ and the distance between them is h. Let $\overline{\mathcal{S}}$ be a square grid imposed on $\mathcal{S}$ with resolution δ_d, which is shown in Fig 7. The same grid $\overline{\mathcal{G}}$ is also imposed on the ground plane $\mathcal{G}$. To demonstrate the main strategy at a high-level, consider a ground point $g \in \mathcal{G}$, where we can find the optimal pair of cameras in the corse camera grid $\overline{\mathcal{S}}$. By moving the ground point in a region $\mathcal{R}(g)$ and using the same pair of cameras, we establish an upper bound of any points in the region $\mathcal{R}(g)$.

In this section, we make this rough argument concrete. Specifically, we show that with the choice of $\delta_d = \frac{h}{\tan(\pi/4-\alpha)}$, we achieve a 2 approximation ratio for uncertainty everywhere using $|\overline{\mathcal{S}}| = \lceil \frac{9Area(\mathcal{G})}{h^2} \rceil$ cameras. Due to space limitations, we present the lemmas and main ideas and leave detailed proofs to the supplementary material.

5.1. Problem 2 in 2D

For cameras in the grid $s \in \overline{S}$ and target $g \in G$, we define the grid uncertainty $\overline{\varepsilon}(g)$ using only the best two cameras in grid $\overline{S}$ as the following

$$\overline{\varepsilon}(g) = \min_{s_i, s_j \in \overline{S}} \varepsilon(g, \{s_i, s_j\})$$

As mentioned earlier, we will choose the grid resolution to be $\delta_d = \frac{h}{\tan(\pi/4-\alpha)}$ which is half of the distance between the optimal pair of cameras.

Lemma 5.1. *The minimum worst case uncertainty of any point g is given as ε_∞, we claim that $\forall g \in \overline{G}$*

$$\overline{\varepsilon}(g) \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} \cdot \varepsilon_\infty$$

In order to bound the uncertainty of any target $\forall g \in G$ using the camera grid $\overline{S}$, we need to explore the uncertainty of the targets in between the grid. Therefore, we fix a grid point and define a range of targets $R(g) = [g - \delta_d/2, g + \delta_d/2]$ such that $R(g)$ is generated by moving $g \in \overline{G}$ along the x-axis of the grid as shown in Fig 9. We now show that the worst case uncertainty is achieved at the end points of this interval (i.e. the midpoint of two grid locations). We do this in two parts: first, we show that the sum of the viewing angles increases as we move away from a grid point until we reach to the midpoint. Next, we use this result to upper bound the uncertainty.

Lemma 5.2. *Let $g \in \overline{G}$ be a grid location with grid uncertainty $\overline{\varepsilon}(g)$ using the optimal pair of cameras $\{s_i, s_j\}$ with orientation θ_i, θ_j respectively. For any $g' \in R(g) = [g - \delta_d/2, g + \delta_d/2]$, s_i and s_j view g' at an angle at most $\theta_i + \theta_j$ is maximized. This value is achieved when the inner half-plane of $Cone_i$ and $Cone_j$ intersects $g \pm \delta_d/2$.*

It turns out that computing the uncertainty directly, as we did for a single point is tricky because the change in the viewing angles is monotonic. Instead we show that $\theta_i + \theta_j$ is monotonic and further, there is a function $\xi(g)$ which increase monotonically with $\theta_i + \theta_j$ and upper bounds $\overline{\varepsilon}(g)$ – i.e $.\overline{\varepsilon}(g) \leq \xi(g)$ for all g. This yields our main result:

Theorem 5.3. *For all targets $g \in G$, the worst case grid uncertainty is bounded as follows*

$$\overline{\varepsilon}(g) \leqslant 1.75\varepsilon_\infty$$

5.2. Problem 2 in 3D

In 3D, we use the same grid resolution $\delta_d = \frac{h}{\tan(\pi/4-\alpha)}$ for the this section, which is half of the distance between the optimal pair of cameras. The main result is

Theorem 5.4. *For all targets $g \in \mathcal{G}$, the worst case grid uncertainty is bounded as follows*

$$\bar{\varepsilon}(g) \leqslant 2\varepsilon_\infty$$

The proof is similar to the 2D case. It is extended to include perturbations in both x and y directions which slightly increases the bounds. See Figures 8 and 9.

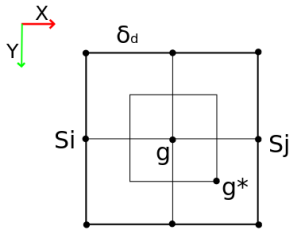

Figure 8: Camera grid: top down view

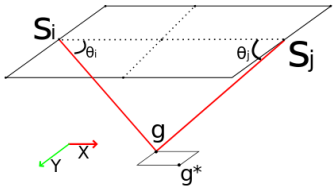

Figure 9: Camera grid in 3D: g is perturbed to g^* to achieve maximum uncertainty.

6. Validation

In this section, we present simulation results used for validating our model and results followed by a real experiment which demonstrates the effectiveness of sensor selection in an aerial photography application:

6.1. Simulation

We used the parameters of a GOPRO HERO 3 for simulations. These were: Resolution: 1920×1080, Field of view: $120^\circ \times 70^\circ$. The calibration error in pixels was $[0.2061, 0.2183]$. For all simulations we used an iMac with 3.3GHz quad-core Intel Core i5 and 16GB of RAM.

Model justification: We consider the following sources of uncertainty: finite resolution, calibration errors, camera center location, camera orientation.

The first two are less than one pixel. To investigate the role of the camera orientation, we perturbed camera location $\hat{s} = s + n_s$, where s is the true location and n_s is a uniform noise, and camera pose $\hat{\theta} = \theta + n_\theta$ where θ is the true orientation and n_θ is a uniform noise. Figure 10 reports the result of triangulation error from two cameras in an optimal position. The height of the viewing plane was set to 100 m. Each simulation was repeated 10^5 times where the target location $\hat{g}$ is computed by triangulation and the error $|\hat{g} - g|$ is reported. Various noise levels are shown in the captions. If we choose a bound of 10 degrees for the measurement error, it corresponds to $\alpha = \max(\frac{10}{1920} * 120^\circ * \frac{\pi}{180}, \frac{10}{1080} * 70^\circ * \frac{\pi}{180}) = 0.0113$ in radians. The solid red line shows the predicted worst case error using our model. The histogram shows that the distance to the true target location is well bounded by the worst case uncertainty which is indicated as the vertical red line.

Next, we study the effect of using two cameras vs. all cameras. We estimate the target pose using least squares from all cameras and report the ratio: $\frac{|\hat{g}_2 - g|}{|\hat{g}_\infty - g|}$ is plotted in Fig 12. Here, $\hat{g}_2$ is the estimated target location using the optimal pair while $\hat{g}_\infty$ uses all the cameras. The noise is set to $|n_p| \leq 10$, $|n_s| \leq 0.003h$, and $|n_\theta| \leq 0.5^\circ$. The simulation was repeated 10^4 times.

The ratio is less than 3.5, which means that using the optimal pair of cameras to triangulate the target is at most 3.5 times worse than triangulation using all cameras.

6.2. Real Experiment

We took video footage by flying low over an area with about 10 meter elevation variation. The flying altitude was set to 20 meters. We recorded around 5 minutes of videos, which is roughly 10000 frames. In order to speed up the reconstruction, we extracted every 25 frames of the videos, which results in 414 frames for the original reconstruction. We use this footage to demonstrate two applications of the proposed grid based sensor selection scheme.

We used the commercial AgiSoft software for Structure from Motion which performs joint error minimization over all pairs of views to investigate the effect of view selection on reconstruction speed and reprojection error. We used 414 frames from the original video and reconstructed the scene. Then, we used three different grids with resolutions $\delta_d = h/3$, $\delta_d = h/2$, and $\delta_d = h$, as shown in Figure 13. The number of frames reduced from 414 to 116, 59, and 20 frames. We re-ran the SFM routine for each set. The total time required to reconstruct the same decreased significantly as shown in Table 6.2. The reprojection error, shown in Fig 14, on the other hand was not significantly impacted.

Table 1: Reconstruction time vs different number of input frames

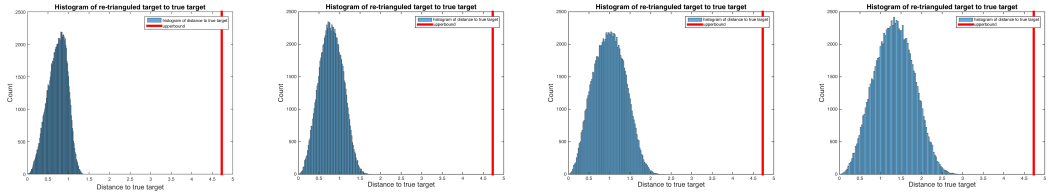

Figure 10: Histogram of the error: $|\hat{g} - g|$ with the following noise parameters. Left to right: (a) $|n_p| \leq 10, |n_s| = 0, |n_\theta| = 0$ (b) $|n_p| \leq 10$, $|n_s| \leq 0.001h$, $|n_\theta| \leq 0.25^\circ$ (c) $|n_p| \leq 10$, $|n_s| \leq 0.003h$, $|n_\theta| \leq 0.5^\circ$ (d) $|n_p| \leq 10, |n_s| \leq 0.006h$, $|n_\theta| \leq 0.75^\circ$

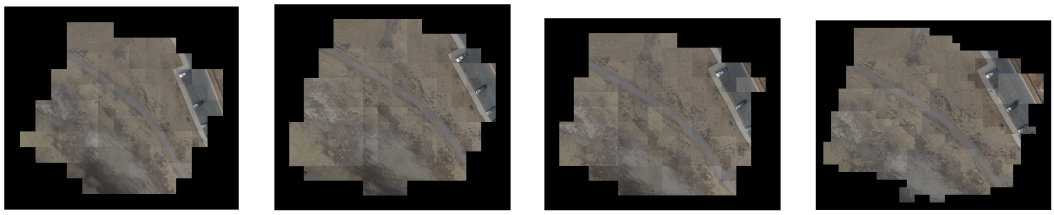

Figure 11: Mosaics from 414, 116, 59 and 20 views

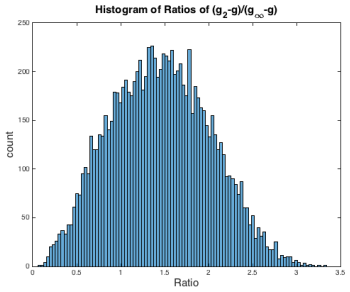

Figure 12: Distribution of $\frac{|\hat{g}_2 - g|}{|\hat{g}_\infty - g|}$

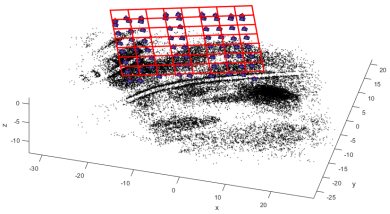

Figure 13: A coarse grid was imposed to select 59 views from the images shown in Figure 1.

Num of Frames	Reconstruction Time(mins)
414	5435
116	1781
59	21
20	2

For qualitative evaluation, we stitched the images together by using the output pose from SFM and orthorectifying the views. to compare the quality of the final mosaic. The results shown in Fig 11 are comparable, indicating that the proposed view selection mechanism does indeed perform comparable with respect to the original input set.

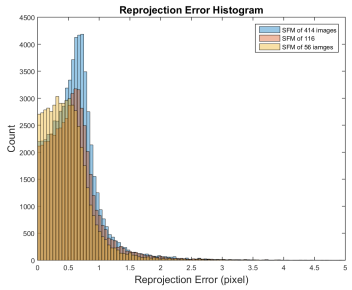

Figure 14: Reprojection error of the reconstructed scene with frames extracted using two different grid resolutions

7. Conclusion

In this paper, we studied view selection for a specific but common setting where a ground plane is viewed from above from a parallel viewing plane. We showed that for a specific world point, two views can be chosen so as to guarantee a reconstruction quality which is almost as good as one that can be obtained by using all possible views. Next, by fixing these two views and studying perturbations of the world point, we showed that one can put a coarse grid on the viewing plane and ensure good reconstructions everywhere. Even though the reconstruction quality can be improved by increasing the grid resolution, we showed that a grid reso-

lution proportional to the scene depth suffices to guarantee a constant factor deviation from the optimal reconstruction.

We also presented an application of these results to image mosaicking from aerial imagery. Our results provide a foundation for multiple avenues of future research. An immediate extension is for scenes which can be represented as surfaces composed of multiple planes. However, as the scene geometry gets more sophisticated, occlusions must be addressed. This raises "art gallery" type research problems [10]. Furthermore, rather than selecting views apriori and in one shot, the view selection can be informed by the reconstruction process as is commonly done in existing literature. Our results can provide a starting point for an iterative scheme where a coarse grid is used for reconstruction under the planar scene assumption and further refined based on the intermediate reconstruction.

References

[1] L. Cheong, C. Fermüller, and Y. Aloimonos. Effects of errors in the viewing geometry on shape estimation. *Computer Vision and Image Understanding*, 71(3):356–372, 1998. 2

[2] H. Farid, S. Lee, and R. Bajcsy. View selection strategies for multi-view, wide-base stereo. Technical report, Technical Report MS-CIS-94-18, University of Pennsylvania, 1994. 2

[3] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Towards internet-scale multi-view stereo. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1434–1441. IEEE, 2010. 2

[4] Y. Gao, M. Wang, Z.-J. Zha, Q. Tian, Q. Dai, and N. Zhang. Less is more: efficient 3-d object retrieval with query view selection. *IEEE Transactions on Multimedia*, 13(5):1007–1018, 2011. 2

[5] A. Hornung, B. Zeng, and L. Kobbelt. Image selection for improved multi-view stereo. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008. 2

[6] V. Isler and M. Magdon-Ismail. Sensor selection in arbitrary dimensions. *IEEE Transactions on Automation Science and Engineering*, 5(4):651–660, 2008. 2

[7] K. N. Kutulakos and C. R. Dyer. Recovering shape by purposive viewpoint adjustment. *International Journal of Computer Vision*, 12(2-3):113–136, 1994. 2

[8] M. Mauro, H. Riemenschneider, A. Signoroni, R. Leonardi, and L. Van Gool. An integer linear programming model for view selection on overlapping camera clusters. In *3D Vision (3DV), 2014 2nd International Conference on*, volume 1, pages 464–471. IEEE, 2014. 2

[9] J. Maver and R. Bajcsy. Occlusions as a guide for planning the next view. *IEEE transactions on pattern analysis and machine intelligence*, 15(5):417–433, 1993. 2

[10] J. O'rourke. *Art gallery theorems and algorithms*, volume 57. Oxford University Press Oxford, 1987. 9

[11] H. Sahabi and A. Basu. Analysis of error in depth perception with vergence and spatially varying sensing. *Computer Vision and Image Understanding*, 63(3):447–461, 1996. 2

[12] W. R. Scott, G. Roth, and J.-F. Rivest. View planning for automated three-dimensional object reconstruction and inspection. *ACM Computing Surveys (CSUR)*, 35(1):64–96, 2003. 2

[13] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Automatic view selection using viewpoint entropy and its application to image-based modelling. In *Computer Graphics Forum*, volume 22, pages 689–700. Wiley Online Library, 2003. 2

Appendix

A.

The proof for Lemma 4.3 is shown here.

Proof. Using small angle approximation, we get $\sin(\alpha) \approx \alpha$ and $\cos(\alpha) \approx 1$ and $\alpha^2 \approx 0$. The angles are constraint such that $\theta_p, \theta_q \in [\pi/4 - 2\alpha, \pi/4]$.

$$\begin{aligned} r_1 &= \frac{t\sin(\theta_q - \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q - 2\alpha)\sin(\theta_p + \theta_q)} \\ &\approx \frac{2t\alpha\sin(\theta_q)}{\sin^2(\theta_p + \theta_q) - \alpha\sin(2\theta_p + 2\theta_q)} \\ &\approx \sqrt{2}t\alpha \end{aligned}$$

$$\begin{aligned} r_2 &= \frac{t\sin(\theta_p + \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q)\sin(\theta_p + \theta_q + 2\alpha)} \\ &\approx \frac{2t\alpha\sin(\theta_p)}{\sin^2(\theta_p + \theta_q) + \alpha\sin(2\theta_p + 2\theta_q)} \\ &\approx \sqrt{2}t\alpha \end{aligned}$$

$$\begin{aligned} r_3 &= \frac{t\sin(\theta_q + \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q)\sin(\theta_p + \theta_q + 2\alpha)} \\ &\approx \frac{2t\alpha\sin(\theta_q)}{\sin^2(\theta_p + \theta_q) + \alpha\sin(2\theta_p + 2\theta_q)} \\ &\approx \sqrt{2}t\alpha \end{aligned}$$

$$\begin{aligned} r_4 &= \frac{t\sin(\theta_p - \alpha)\sin(2\alpha)}{\sin(\theta_p + \theta_q - 2\alpha)\sin(\theta_p + \theta_2)} \\ &\approx \frac{2t\alpha\sin(\theta_p)}{\sin^2(\theta_p + \theta_q) - \alpha\sin(2\theta_p + 2\theta_q)} \\ &\approx \sqrt{2}t\alpha \end{aligned}$$

$$\begin{aligned} diag_1 &= ||r_1^2 + r_2^2 - 2 \cdot r_1 \cdot r_2 \cdot \cos(\theta_p + \theta_q)||_2 \\ &\approx ||2t^2\alpha^2 + 2t^2\alpha^2 - 4t^2\alpha^2\cos(\theta_p + \theta_q)||_2 \\ &= 2t\alpha||1 - \cos(\theta_p + \theta_q)||_2 \end{aligned}$$

$\max(diag_1) \leq 2t\alpha$ and $\min(diag_1) \geq \sqrt{1 - 4\alpha} \cdot 2t\alpha$

$$\begin{aligned} diag_2 &= ||r_1^2 + r_4^2 - 2 \cdot r_1 \cdot r_4 \cdot \cos(\pi - \theta_p - \theta_q + 2\alpha)||_2 \\ &\approx 2t\alpha||1 + \cos(\theta_p + \theta_q) - 2\alpha\sin(\theta_p + \theta_q)||_2 \end{aligned}$$

$\max(diag_2) \leq \sqrt{1 + 2\alpha} \cdot 2t\alpha$ and $\min(diag_2) \geq \sqrt{1 - 4\alpha} \cdot 2t\alpha$. Therefore, $diag_2 \leq \sqrt{\frac{1+2\alpha}{1-4\alpha}} diag_1$ and $1 - 4\alpha$ will not be negative since α must be less than 0.25 to satisfy small angle approximation.

Given that $\varepsilon_2 = \max(diag1, diag2)$, we can conclude that

$$\begin{aligned} \varepsilon_2 &\leq \frac{1 + 2\alpha}{1 - 4\alpha}\frac{2h\sin(2\alpha)}{1 - \sin(2\alpha)} \\ &\lesssim \frac{1 + 2\alpha}{1 - 4\alpha}\frac{4h\alpha}{1 - 2\alpha} \end{aligned}$$

□

B.

The proof for Lemma 5.1 is shown here.

Proof. Since the grid uncertainty only uses cameras from the grid $s \in \overline{S}$ and the δ_d is half of the distance between the optimal pair of cameras, $\forall g \in \overline{G}$, $\exists\{s_i, s_j\} \subseteq \overline{S}$ such that $\{s_i, s_j\}$ is the optimal pair for g. Therefore, the minimum grid uncertainty for all target in $g \in \overline{G}$ behaves the same as in Theorem 4.1. □

C.

The proof of Lemma 5.2 is shown here.

Proof. First, it is clear that both θ_i and θ_j are increased when the inner half-plane of the both cones intersects the target. Now assume target g is moved to along the x axis as shown in Fig 9 by some length m, where $m \leq \delta_d/2$. We can formulate $\theta_i + \theta_j$ as a function of m and the distance between the cameras as

$$\begin{aligned} f(m) = \theta_i + \theta_j = \tan^{-1}\left(\frac{h}{h/\tan(\pi/4 - \alpha) - m}\right) + \\ \tan^{-1}\left(\frac{h}{h/\tan(\pi/4 - \alpha) + m}\right) + 2\alpha \end{aligned}$$

The derivative of $f(x)$ is given as

$$\begin{aligned} \frac{d}{dm}f(m) = \{2m(2\cos(2\alpha) + 2\cos(2\alpha)\sin(2\alpha)\} \cdot \\ \{2m^2\sin(2\alpha) + 2m^4\sin(2\alpha) + 4m^2\sin^2(2\alpha) \\ + m^4\sin^2(2\alpha) + m^4 + 4\}^{-1} \end{aligned}$$

It is clear that the derivative is always positive. Therefore, $\theta_i + \theta_j$ is a monotonically increasing function with m, which concludes the proof. □

D.

We show the proof for Theorem 5.3.

Proof. Rather than computing the uncertainty directly, we modify the diagonal equations Eq 7 and Eq 8 such that $\bar{\varepsilon}(g) \leq \xi(g)$ with the same pair of cameras and $\xi(g)$ increase monotonically when $\theta_i + \theta_j$ increases. We can use $\xi(g)$ to establish an upper bound for $\bar{\varepsilon}(g)$.

To do this, we modified the diagonal Eq 7 and Eq 8 as

$$r'_1 = r'_4 = t\frac{\sin(\max(\theta_i,\theta_j)-\alpha)\sin(2\alpha)}{\sin(\theta_i+\theta_j-2\alpha)\sin(\theta_i+\theta_j)} \tag{11}$$

$$r'_2 = r'_3 = t\frac{\sin(\max(\theta_i,\theta_j)+\alpha)\sin(2\alpha)}{\sin(\theta_i+\theta_j)\sin(\theta_i+\theta_j+2\alpha)} \tag{12}$$

where $t = \frac{2h}{\tan(\pi/4-\alpha)}$,

$$diag'_1 = ||r'^2_1 + r'^2_2 - 2r'_1r'_2\cos(\theta_i+\theta_j)||_2 \tag{13}$$

$$diag'_2 = ||r'^2_1 + r'^2_4 + 2r'_1r'_4\cos(2\min(\theta_i,\theta_j) - 2\alpha)||_2 \tag{14}$$

We now define the modified uncertainty function as

$$\xi(g,\{s_i,s_j\}) = \max(diag'_1, diag'_2, r'_1, r'_2, r'_3, r'_4) \tag{15}$$

The key property of ξ is monotonicity w.r.t. $\theta_i + \theta_j$. Therefore, using the function ξ and Lemma 5.2, we can conclude that $\xi(g + x_a, \{s_i, s_j\}) \leq \xi(g + x_b, \{s_i, s_j\})$ for $0 \leq x_a \leq x_b \leq \delta_d/2$.

Using small angle approximation, we get $\sin(\alpha) \approx \alpha$ and $\cos(\alpha) \approx 1$ and $\alpha^2 \approx 0$. Given the orientation at the location $g + \delta_d$ as $\theta_i = \tan^{-1}(\frac{2}{3}\tan(\pi/4-\alpha)) + \alpha$ and $\theta_j = \tan^{-1}(2\tan(\pi/4-\alpha)) + \alpha$, we get the following:

$$\begin{aligned}\theta_i &= \tan^{-1}(\frac{2}{3}\tan(\pi/4-\alpha)) + \alpha \\ &\leq \tan^{-1}(2/3) + \alpha \\ &\leq 0.589 + \alpha\end{aligned}$$

$$\begin{aligned}\theta_j &= \tan^{-1}(2\tan(\pi/4-\alpha)) + \alpha \\ &\leq \tan^{-1}(2) + \alpha \\ &\leq 1.1072 + \alpha\end{aligned}$$

while $\theta_i + \theta_j \leq 1.6952 + 2\alpha$. Therefore, for Eq 13 and Eq 14, we can derive the following equations.

$$\begin{aligned}r'_1 = r'_4 &= t\frac{\sin(\max(\theta_i,\theta_j)-\alpha)\sin(2\alpha)}{\sin(\theta_i+\theta_j-2\alpha)\sin(\theta_i+\theta_j)} \\ &\lesssim \frac{2t\sin(1.1072)\alpha}{\sin(1.6952)\sin(1.6952+2\alpha)}\end{aligned}$$

Since $\alpha \leq 0.1$,

$$\leq 1.91t\alpha$$

$$\begin{aligned}r'_2 = r'_3 &= t\frac{\sin(\max(\theta_i,\theta_j)+\alpha)\sin(2\alpha)}{\sin(\theta_i+\theta_j+4\alpha)\sin(\theta_i+\theta_j+2\alpha)} \\ &\lesssim \frac{2t\sin(1.1072)\alpha}{\sin(1.6952+4\alpha)\sin(1.6952+2\alpha)}\end{aligned}$$

Since $\alpha \leq 0.1$,

$$\leq 2.19t\alpha$$

$$\begin{aligned}diag'_1 &= ||r'^2_1 + r'^2_2 - 2r'_1r'_2\cos(\theta_i+\theta_j)||_2 \\ &\leq t\alpha||1.91^2 + 2.19^2 - 8.37\cos(1.6952+2\alpha)||_2\end{aligned}$$

Since $\alpha \leq 0.1$,

$$\leq 3.34t\alpha$$

$$\begin{aligned}diag'_2 &= ||r'^2_1 + r'^2_4 + 2r'_1r'_4\cos(2\theta_i - 2\alpha)||_2 \\ &\leq t\alpha||1.91^2 + 2.19^2 + 8.37\cos(0.589*2)||_2 \\ &\leq 3.5t\alpha\end{aligned}$$

Given $2t\alpha \leq \varepsilon_\infty$, we can bound

$$\begin{aligned}\xi(g+\delta_d) &= diag'_2 \\ &\leq 3.5t\alpha \\ &\leq \frac{3.5t\alpha}{2t\alpha}\cdot\varepsilon_\infty \\ &\leq 1.75\varepsilon_\infty\end{aligned}$$

□

E.

Theorem 5.4 is proved here.

Proof. We can calculate the angle at location g'.

$$\begin{aligned}\theta_i &= \tan^{-1}\{\frac{2}{3}\frac{(h^2+[h/(2\tan(\pi/4-\alpha))]^2)^{0.5}}{h/\tan(\pi/4-\alpha)}\} + \alpha \\ &= \tan^{-1}\{\frac{2}{3}\sqrt{\tan^2(\pi/4-\alpha)+1/4}\} + \alpha \\ &\leq \tan^{-1}(\frac{2}{3}\sqrt{1.25}) + \alpha \\ &\leq 0.641 + \alpha\end{aligned}$$

$$\begin{aligned}\theta_i &= \tan^{-1}\{2\frac{(h^2+[h/(2\tan(\pi/4-\alpha))]^2)^{0.5}}{h/\tan(\pi/4-\alpha)}\} + \alpha \\ &= \tan^{-1}\{2\sqrt{\tan^2(\pi/4-\alpha)+1/4}\} + \alpha \\ &\leq \tan^{-1}(2\sqrt{1.25}) + \alpha \\ &\leq 1.151 + \alpha\end{aligned}$$

Therefore, $\theta_i + \theta_j \leq 1.792 + 2\alpha$.

$$\begin{aligned} r'_1 = r'_4 &= t\frac{\sin(\max(\theta_i, \theta_j) - \alpha)\sin(2\alpha)}{\sin(\theta_i + \theta_j - 2\alpha)\sin(\theta_i + \theta_j)} \\ &\lesssim \frac{2t\sin(1.151)\alpha}{\sin(1.792)\sin(1.792 + 2\alpha)} \\ Since\ \alpha \leq 0.1, \\ &\leq 2.051t\alpha \end{aligned}$$

$$\begin{aligned} r'_2 = r'_3 &= t\frac{\sin(\max(\theta_i, \theta_j) + \alpha)\sin(2\alpha)}{\sin(\theta_i + \theta_j)\sin(\theta_i + \theta_j + 2\alpha)} \\ &\lesssim \frac{2t\sin(1.151)\alpha}{\sin(1.792 + 4\alpha)\sin(1.792 + 2\alpha)} \\ Since\ \alpha \leq 0.1, \\ &\leq 2.461t\alpha \end{aligned}$$

$$\begin{aligned} diag'_1 &= ||{r'}_1^2 + {r'}_2^2 - 2r'_1r'_2\cos(\theta_i + \theta_j)||_2 \\ &\leq t\alpha||2.051^2 + 2.461^2 - 10.1\cos(1.792 + 2\alpha)||_2 \\ Since\ \alpha \leq 0.1, \\ &\leq 4t\alpha \end{aligned}$$

$$\begin{aligned} diag'_2 &= ||{r'}_1^2 + {r'}_4^2 + 2r'_1r'_4\cos(2\theta_i - 2\alpha)||_2 \\ &\leq t\alpha||2.051^2 + 2.461^2 + 10.1\cos(0.641 * 2)||_2 \\ &\leq 3.63t\alpha \end{aligned}$$

Given $2t\alpha \leq \varepsilon_\infty$, we can bound

$$\begin{aligned} \xi(g') &= diag'_2 \\ &\leq 4t\alpha \\ &\leq \frac{4t\alpha}{2t\alpha} \cdot \varepsilon_\infty \\ &\leq 2\varepsilon_\infty \end{aligned}$$

□